

# A Survey on Reproducibility in Parallel Computing

Sascha Hunold

Vienna University of Technology  
Faculty of Informatics  
Research Group for Parallel Computing  
Favoritenstraße 16/184-5  
1040 Vienna, Austria  
hunold@par.tuwien.ac.at

## Abstract

We summarize the results of a survey on reproducibility in parallel computing, which was conducted during the Euro-Par conference in August 2015. The survey form was handed out to all participants of the conference and the workshops. The questionnaire, which specifically targeted the parallel computing community, contained questions in four different categories: general questions on reproducibility, the current state of reproducibility, the reproducibility of the participants' own papers, and questions about the participants' familiarity with tools, software, or open-source software licenses used for reproducible research.

## 1 Introduction

Conducting sound and reproducible experiments in parallel computing is not easy, as hardware and software architectures of current parallel computers are most often very complex. This high complexity makes it difficult and often impossible for scientists to model such systems mathematically. Thus, scientists often rely on experiments to study new parallel algorithms, different software solutions (e.g., operating systems), or novel hardware architectures. The situation in parallel computing is made even more difficult as parallel systems are in a constant state of flux, e.g., the total core count is rapidly growing and many programming paradigms for parallel machines have emerged and are actively being used.

We established the first edition of the International Workshop on Reproducibility in Parallel Computing (REPPAR<sup>1</sup>) in conjunction with the Euro-Par conference in 2014. The workshop is concerned with experimental practices in parallel computing research. It should be a forum for discussing and exchanging ideas to improve reproducibility matters in our research domain. We solicit research papers and experience reports on a number of relevant topics, particularly: methods for analysis and visualization of experimental data, best practice recommendations, results of attempts to replicate previously published experiments, and tools for experimental computational sciences. Some examples of the latter include workflow management systems, experimental test-beds, and systems for archiving and querying large data files.

In 2015, the REPPAR workshop was hosted for the second time in conjunction with the Euro-Par conference. This year we wanted to spark a fruitful discussion by conducting a survey on reproducible research and by evaluating the results directly during the workshop. In the present paper, we will take a closer look at the results of the survey and discuss some of the findings.

After summarizing related work in Section 2, we explain the context of the survey in Section 3. Section 4 presents the survey results, and we draw conclusions in Section 5.

---

<sup>1</sup><http://reppar.org/>

## 2 Related Work

Improving the reproducibility of results that get published in today’s scientific journals is one of the big challenges of the current research landscape, not only because the problem has lately been brought into the spotlight by journals like *Science* or *Nature* (cf. [1, 2]). Thus, many researchers across disciplines are trying to tackle the problem of the irreproducibility of scientific findings.

From a computer-science standpoint, we are foremost interested in the state of reproducibility of computational results. The reproduction of scientific findings in computational sciences has other challenges than, say, medicine, as here we study abstract objects, e.g., a computer program or an algorithm (rather than the human body). Questions that arise in this context are, for example, how to share source code (technically) or which license to apply to a piece of software? Stodden, Leisch, and Peng addressed these issues and published a collection of articles, in which several solutions to the dilemma are proposed [7].

Here, we are not only interested in computational sciences, but specifically in parallel and distributed computing, where we are facing additional challenges in terms of reproducibility [3]. For example, in the high performance computing community, scientists are primarily interested in optimizing performance, e.g., trying to minimize the run-time or to maximize the throughput of a system. Thus, a reproducible analysis does not only need to be able to solve the computational problem with the same outcome, but also in the same—or at least comparable—time as shown in the original paper.

Therefore, we conducted a survey on reproducible research among the Euro-Par participants to gain insights about how the reproducibility problem is perceived in our community. In the USA, several initiatives or workshops exist that address the reproducibility problem for large-scale computing. One example is the XSEDE workshop on reproducibility [4].

Several surveys have been undertaken in the broader context of reproducible research. For us, most related to our work are the survey of Stodden [6] and the survey of Prabhu et al. [5]. Stodden’s survey sheds light on the incentives for scientists to share or not to share their work (code or data). The survey by Prabhu et al. is more concerned with best practices in computational sciences, for example, the authors try to answer questions like “do scientists know about parallelization techniques for speeding up their applications” or “what languages do scientists use for their daily compute tasks”. The survey results by Prabhu et al. reveal that “[s]cientists should release code to their peers” in order to “allow other scientists to reproduce prior work” [5].

## 3 Context of the Survey

To conduct our survey on reproducibility in parallel computing, we prepared a questionnaire containing 24 questions (cf. Appendix A), which we grouped into four different categories (cf. Section 4.1–Section 4.4). All participants of the Euro-Par conference received one survey flyer, which advertised and introduced the survey to them. The flyer contained a unique token that enabled each participant to vote exactly once. The survey was completely anonymous and since flyers (and their tokens) were handed out in the order in which participants arrived at the conference registration desk, the identity of the voters was additionally protected. The survey was implemented using the LimeSurvey software<sup>2</sup>. We printed 300 survey flyers, each containing one token, and handed out one flyer to each of the 232 participants of the Euro-Par conference. Unfortunately, only 31 persons (13%) completed the online questionnaire.

---

<sup>2</sup><http://www.limesurvey.org/>

## 4 Survey Results

Now, we present the survey results and comment on the outcome of individual questions.


### 4.1 General Questions on Reproducibility

The first question (Q2.1.1) directly asked whether the survey participant is interested in reproducible research. Rather surprisingly, the majority of the participants ( $> 90\%$ ) declared their interest in reproducibility. Considering the fact that only 31 persons completed the survey, we conclude that most of the participants also attended the REPPAR workshop. Therefore, we should keep in mind that our results are highly biased towards a small group of people sharing similar interests.


We also assumed that only few scientists know what “reproducible research” means and also what the difference is between the terms “replicability”, “repeatability”, and “reproducibility”. Our assumption was based on the fact that many articles use different definitions of “reproducibility”. Since we had posed a vague question, the poll results of questions Q2.1.2 and Q2.1.3 were surprising. For example, only 13% of the voters stated that they do not know the difference between replicability, repeatability, and reproducibility.

It is also noteworthy that all survey participants think that the reproduction of already published results is worth another publication. However, 65% of them demand that articles reproducing work of others need to contain new insights.


#### Q4.1.1 Do you care (in general) about the reproducibility of scientific results (your own, others)?




#### Q4.1.2 Do you know what people mean when speaking about "reproducible" results?




#### Q4.1.3 Do you know the difference between replicability, repeatability, and reproducibility?



#### Q4.1.4 Do you think that the reproduction of already published results is worth another publication?



#### Q4.1.5 Have you tried to reproduce the results of others?




## 4.2 Current State of Reproducibility in Parallel Computing

With the second block of questions, we intended to learn more about how scientists see the domain of parallel or high performance computing (HPC) in terms of reproducibility.


The results of question Q2.2.1 show a clear picture: almost all participants think that the reproducibility of articles in our research domain need to be improved. We again note that our results are biased, as many of the survey participants also attended the REPPAR workshop on reproducible research.

It is also remarkable that only a small percentage (6%) of the voters believed that articles from top conferences such as PPoPP or IPDPS are better reproducible than papers of other conferences (Q2.2.3). On the contrary, many people (>50% in sum) do not trust the reproducibility of the results when they review scientific articles (Q2.2.4).


**Q4.2.1 Do you think the state of reproducibility for articles in our research domain (Parallel Computing/HPC) needs to be improved?**




**Q4.2.2 Do you think current research articles in the domain of Parallel Computing/HPC are reproducible by other independent researchers?**



**Q4.2.3 Do you think that results published in top conferences (e.g., PPOPP, IPDPS) are generally easier to reproduce than those published in lower-tier conferences in parallel computing (in the last 5 years)?**



#### Q4.2.4 How often do you question the reproducibility of results when you review other scientific articles?



### 4.3 Reproducibility of Your Articles

The third block of questions was concerned with what the participants think about the reproducibility of their own articles. The poll results for question Q2.3.1 show that a significant fraction of the voters (19%) believe that the results published in their articles are reproducible by others. Surprisingly, only 3% stated that they know that their papers are not reproducible. We had expected a higher percentage of people that would admit that their papers are hard to reproduce, especially when taking into account that the poll was anonymous.


90% of the participants consider freely accessible HPC systems a necessity for reproducible results.

We also asked how scientists provide the source code, the raw experimental data, and the data analysis procedures to others. Again, it was surprising that a large percentage (23%) of scientists stated that they publish the source code along with their papers (Q2.3.2). From our personal experience we had expected much less (around 10%). The poll results also show that more than half of the scientists use a public revision control system, such as GitHub, to share their code (Q2.3.4).


However, when we look at the percentage of scientists that do not provide the source code, the raw experimental data, or the data analysis procedures, we can observe that the data analysis procedures get shared less often compared to the other two. One explanation could be that the data analysis procedures applied are very simple (e.g., computing the arithmetic mean). Another explanation could be that researchers simply do not give them a high priority, and perhaps do not see the importance for others to have these procedures.

We also asked the survey participants about their main reasons for not sharing code, data, or data analysis procedures (Q2.3.8). Here, no clear line can be drawn, as no answer was mentioned significantly more often than others. Similarly, we did not obtain a clear picture when asking the participants what they believe are the major obstacles to reproduce their papers (Q2.3.9).


#### Q4.3.1 Do you think the results (contribution) published in YOUR papers are reproducible by others (in the last 5 years)?




**Q4.3.2 How often have you published the source code along with YOUR paper (in the last 5 years)?**




**Q4.3.3 Do you consider freely accessible HPC systems a necessity for getting reproducible performance figures?**




**Q4.3.4 How do you provide "source code" for others?**




**Q4.3.5 How do you provide the "raw data (of experiments)" for others?**




**Q4.3.6 How do you provide the "data analysis procedure (R scripts, etc)" for others?**



**Q4.3.7 How do you document how to use your source code / data analysis scripts for others?**




**Q4.3.8 What are the main reasons for NOT making the source code/raw data/data analysis procedure available?**





**Q4.3.9 What do you think will be the main difficulties/obstacles when other independent researchers try to reproduce YOUR experiments?**




## **4.4 Tools, Software, and Licenses for Reproducible Research**

Last, we wanted to examine which software and licenses scientists use for making their experiments reproducible.


In question Q2.4.1, we asked the participants whether they use statistical software packages, such as R or SPSS, for performing data analysis tasks. It turns out that only a third of the voters use such tools on a regular basis. It is also remarkable that most of the voters (71% and 84% respectively) had never used software for literate programming (e.g., knitr or org-mode) nor tools for managing or executing scientific workflows (e.g., VisTrails or Kepler).

Researchers often debate what open-source software license is the best for their purposes. We therefore asked the question whether the participants do know the license policy of their research institutions (Q2.4.6). Only 19% of the voters know this policy, whereas 26% stated that the institute has no explicit policy.


**Q4.4.1 Have you used statistical software packages (e.g., R, SAS, SPSS) for analyzing your experimental results?**




**Q4.4.2 How would you rate YOUR knowledge of the programming language "R"?**




**Q4.4.3 Do you use/have you used tools for literate programming (e.g., knitr, org-mode) for publishing articles?**




**Q4.4.4 Do you have practical experiences with workflow tools to support reproducible research (e.g., VisTrails, Kepler, DataMill, etc.)?**



**Q4.4.5 Do you know the differences between the available common open source licenses?**



#### Q4.4.6 Do you know the policy used by YOUR research institution concerning the choice of open source licenses?



## 5 Conclusions

We presented the poll results of a survey on reproducible research, which had been conducted during the Euro-Par conference 2015. Despite the fact that only 31 persons completed the survey, the results give us some evidence that reproducibility is a problem in our domain. In fact, the survey revealed that the majority of the voters believe that the state of reproducibility needs to be improved in the domain of parallel and high performance computing. The survey participants also think that the majority of the results presented in papers that they receive for review are unlikely to be reproducible. The survey also showed that scientists need to be better informed what the different open-source licenses actually mean and which licenses are allowed to be applied by their research institutions. Last, we found evidence that many scientists are not familiar with software for literate programming and with scientific workflows, which can potentially help to improve reproducibility of articles.

## References

- [1] Announcement: Reducing our irreproducibility. *Nature*, 496:398–398, 2013.
- [2] S. Buck. Solving reproducibility. *Science*, 348(6242):1403, 2015.
- [3] S. Hunold and J. L. Träff. On the state and importance of reproducible experimental research in parallel computing. *CoRR*, abs/1308.3648, 2013.
- [4] D. James, N. Wilkins-Diehr, V. Stodden, et al. Standing together for reproducibility in large-scale computing: Report on reproducibility@XSEDE. *CoRR*, abs/1412.5557, 2014.
- [5] P. Prabhu, T. B. Jablin, A. Raman, Y. Zhang, J. Huang, H. Kim, N. P. Johnson, F. Liu, S. Ghosh, S. Beard, T. Oh, M. Zoufaly, D. Walker, and D. I. August. A survey of the practice of computational science. In *State of the Practice Reports*, SC '11, pages 19:1–19:12, New York, NY, USA, 2011. ACM.
- [6] V. Stodden. The Scientific Method in Practice: Reproducibility in the Computational Sciences. Technical Report Working Paper 4773-10, MIT Sloan School of Management, Feb. 2010.
- [7] V. Stodden, F. Leisch, and R. D. Peng, editors. *Implementing Reproducible Research*. The R Series. Chapman and Hall/CRC, Apr. 2014.

## A Original Questionnaire

### A.1 General Questions on Reproducibility

1. Do you care (in general) about the reproducibility of scientific results (your own, others)?
  - (a) no
  - (b) yes
2. Do you think you know what people mean when speaking about "reproducible" results?
  - (a) no
  - (b) not sure, but I guess so
  - (c) sure, I know what that means
3. Do you know the difference between replicability, repeatability, and reproducibility?
  - (a) no
  - (b) I am not sure, but I guess so
  - (c) sure, I know what the differences are
4. Do you think that the reproduction of already published results is worth another publication?
  - (a) yes, the reproduction alone is worth another publication
  - (b) yes, but only if the publication contains new insights
  - (c) no
5. Have you tried to reproduce the results of others?
  - (a) no
  - (b) I tried once or twice
  - (c) a couple of times ( $> 2$  and  $\leq 10$ )
  - (d) many times ( $> 10$ )

### A.2 Current State of Reproducibility in Parallel Computing

1. Do you think the state of reproducibility for articles in our research domain (Parallel Computing/HPC) needs to be improved?
  - (a) no
  - (b) yes
2. Do you think current research articles in the domain of parallel computing/HPC are reproducible by other independent researchers?
  - (a) yes, all of them
  - (b) yes, except a few papers (90% reproducibility)
  - (c) 50/50, some are, some are not
  - (d) no, but a few might be reproducible (10% reproducibility)
  - (e) no article is reproducible ( $< 1\%$  reproducibility)
  - (f) I really do not know

3. Do you think that results published in top conferences (e.g., PPOPP, IPDPS) are generally easier to reproduce than those published in lower-tier conferences in parallel computing (in the last 5 years)?
  - (a) yes, I know from my experience
  - (b) probably, I can imagine that
  - (c) I am not sure, but I guess not
  - (d) no, not at all (all equally reproducible or non-reproducible)
4. How often do you question the reproducibility of results when you review other scientific articles?
  - (a) for more than 90% of the articles
  - (b) for more than 50% of the articles
  - (c) for more than 10% of the articles
  - (d) never

### **A.3 Reproducibility of Your Articles**

1. Do you think the results (contribution) published in YOUR papers are reproducible by others (in the last 5 years)?
  - (a) yes, 100%
  - (b) most of them should be (>50%)
  - (c) I am not sure, I guess most results will not be ( $\leq 50\%$  reproducibility)
  - (d) honestly, I know that they are not! ( $<5\%$  reproducibility)
2. How often have you published the source code along with YOUR paper (in the last 5 years)?
  - (a) never
  - (b) very few times ( $<25\%$ )
  - (c)  $\geq 25\%$  and  $< 50\%$
  - (d)  $\geq 50\%$
  - (e) 100% (for each article)
3. Do you consider freely accessible HPC systems a necessity for getting reproducible performance figures?
  - (a) yes, for all studies
  - (b) yes, but only for some studies
  - (c) no
4. How do you provide "source code" for others?
  - (a) I do not provide the source code
  - (b) as an email attachment in response to a direct request
  - (c) as an archive (zip, tar) on a personal webpage
  - (d) I use public revision control system (e.g., GitHub)
5. How do you provide the "raw data (of experiments)" for others?
  - (a) I do not provide the raw data
  - (b) as an email attachment in response to a direct request

- (c) as an archive (zip, tar) on a personal webpage
  - (d) I use public revision control system (e.g., GitHub)
6. How do you provide the "data analysis procedure (R scripts, etc)" for others?
- (a) I do not provide the data analysis procedure
  - (b) as an email attachment in response to a direct request
  - (c) as an archive (zip, tar) on a personal webpage
  - (d) I use public revision control system (e.g., GitHub)
7. How do you document how to use your source code / data analysis scripts for others?
- (a) I do not document them
  - (b) simple README files
  - (c) standard documentation system (e.g., doxygen)
  - (d) electronic laboratory notebook
8. What are the main reasons for NOT making the source code/raw data/data analysis procedure available? (multiple options)
- (a) it does not apply to me (as I make them available)
  - (b) Technical difficulties. Lack of suited tools or hosting infrastructure
  - (c) Institution policy or legal aspects
  - (d) I want to retain a competitive advantage
  - (e) it is too time consuming
  - (f) it is not rewarding
  - (g) it is irrelevant because evolution is too fast
  - (h) other
9. What do you think will be the main difficulties/obstacles when other independent researchers try to reproduce your experiments? (multiple options)
- (a) the lack of access to specific machines
  - (b) the lack of a specific software setup
  - (c) the lack of documentation
  - (d) the lack of time to reproduce our results
  - (e) the lack of scientific credits (others will not get many credits for reproducing our results)
  - (f) other

#### **A.4 Tools/Software/Licenses for Reproducible Research**

1. Have you used statistical software packages (e.g., R, SAS, SPSS, ..) for analyzing your experimental results?
  - (a) no, not at all
  - (b) not on a regular basis
  - (c) yes, I always use them
2. How would you rate YOUR knowledge of the programming language "R"?

- (a) I have never heard of R
  - (b) I am a novice
  - (c) I can code if needed, but I would not call myself an expert
  - (d) I am an advanced user (expert)
3. Do you use/have you used tools for literate programming (e.g., knitr, org-mode, ..) for publishing articles?
- (a) never
  - (b) I used them, but was not convinced
  - (c) I used them and will do so in the future
  - (d) I always use them
4. Do you have practical experiences with workflow tools to support reproducible research (e.g., VisTrails, Kepler, DataMill, etc.)?
- (a) I never used them
  - (b) I used them, but they were not convincing
  - (c) I used them several times and I plan to use them in the future
  - (d) I now use them all the time
5. Do you know the differences between the available common open source licenses?
- (a) no, I have no clue
  - (b) I know the basic differences
  - (c) yes, I have a solid background on licenses
6. Do you know the policy used by YOUR research institution concerning the choice of open source licenses?
- (a) I have no idea
  - (b) not really, but I know where to look/who to ask
  - (c) yes, I know the policy
  - (d) yes, and it is that there is no explicit policy